

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Method and Apparatus for
Selecting Cache and Proxy Policy**

Inventors:

Dawson F. Dean
Chadd B. Knowlton
Mark D. VanAntwerp
Bret P. O'Rourke
Chih-Kan Wang
David J. Roth

ATTORNEY'S DOCKET NO. MS1-575US

1 **TECHNICAL FIELD**

2 The present invention relates to media systems and, more particularly, to
3 the handling of cache and proxy policies in a cache server.

4

5 **BACKGROUND**

6 Media distribution systems distribute media content, such as audio data or
7 video data, from an origin server to one or more clients via a data communication
8 network. The data communication network may include, for example, a local area
9 network (LAN) or a wide area network (WAN). The origin server is the source of
10 the media content, including both live media content and pre-recorded media
11 content. A particular type of media distribution system, referred to as a streaming
12 media system, distributes real-time data from an origin server to one or more
13 clients requesting the real-time data. Increased bandwidth through the data
14 communication network improves the audio or video quality of the distributed
15 media in the case of real-time data. If one or more low bandwidth communication
16 links exist in the network connecting the origin server to the client, then the
17 quality of the real-time media presentation is reduced. In this situation, the client
18 may attempt to retrieve the media from a different server, referred to as a cache
19 server or a proxy server, which is connected to the client using higher bandwidth
20 communication links. The cache or proxy server retrieves the content from the
21 origin server and distributes the content to one or more clients via a faster
22 communication link.

23 Fig. 1 illustrates an environment 100 in which a client can request and
24 receive content (such as streaming video content) from a cache server and/or an
25

1 origin server. A client 102 is coupled to a network 104, such as the Internet, via a
2 communication link 114. A cache server 106 is coupled to client 102 via a
3 communication link 112 and coupled to network 104 via a communication link
4 118. Cache server 106 includes a storage disk 108 for storing data, such as cached
5 media content for distribution to client 102. An origin server 110 is coupled to
6 network 104 via a communication link 116.

7 Typically, the communication links to network 104 (i.e., communication
8 links 114, 116, and 118) are relatively slow connections (e.g., 64k bits per second)
9 and provide low quality real-time video images. In contrast, communication link
10 112 between client 102 and cache server 106 is generally a faster connection (e.g.,
11 100M bits per second) provided by a LAN or other high-speed network.

12 When client 102 wants to receive media content, a request is issued to
13 cache server 106. If the cache server 106 has cached a copy of the requested
14 media content, then the cache server transmits the content across communication
15 link 112 to client 102.

16 Cache server 106 receives its stored media content from one or more origin
17 servers 110. Typically, cache server 106 downloads content from origin server 110
18 at a relatively low speed. However, once the content is downloaded, the cache
19 server 106 can distribute the content to many local clients via a high-speed
20 network connection, such as communication link 112. The cache server 106 may
21 save the downloaded content to serve future client requests for the content without
22 having to download the content from the origin server 110. If cache server 106
23 does not contain the content requested by the client 102, then the client must
24 retrieve the desired content from the origin server 110.
25

1 Existing cache servers are manufactured with a substantially fixed
2 architecture that is difficult to customize. A particular manufacturer may provide
3 several different cache server models with different features, but the architecture
4 of each model provides minimal opportunity for customization. A cache server
5 customer is forced to select a cache server model that is “closest” to their cache
6 requirements. Since different cache users are likely to have different cache
7 requirements, no single cache server can satisfy the needs of all users. Depending
8 on the cache server features and the customer’s requirements, the cache server
9 may not be capable of meeting all of the customer’s cache requirements. These
10 existing cache servers perform various cache-related functions that are controlled
11 by a set of cache policies that are determined by the manufacturer. These existing
12 cache servers that define both the cache-related functions and the policies for
13 implementing those functions minimize the opportunities for a customer to modify
14 the cache server to meet their specific needs.

15 The system described herein addresses these limitations by providing a
16 single cache server that can operate as either an origin server or a cache server.
17 The system includes a flexible architecture that separates the various cache
18 functions from the policies that determine which functions are to be performed in
19 a particular situation. This flexible architecture allows a common cache server to
20 be configured in a variety of different ways depending on the set of policies
21 applied to the cache server. Thus, different manufacturers or end users can create
22 different cache server functionality based on the set of policies selected.

1 **SUMMARY**

2 The system and methods described herein provide a cache server that is
3 capable of performing various cache-related functions under the direction of a set
4 of user-configurable cache policies. This cache server architecture allows the
5 same cache server to operate in different ways depending on the specific set of
6 cache policies.

7 In one embodiment, a media serving engine is provided to distribute media
8 content. A cache engine, which is coupled to the media serving engine, caches
9 media content. A set of cache policies is accessible by the cache engine to define
10 the operation of the cache engine.

11 In a described embodiment, the apparatus can be configured to operate as a
12 cache server.

13 A particular embodiment also includes a data communication interface
14 coupled to the cache engine and the media serving engine. The data
15 communication interface allows the cache engine to retrieve media content across
16 a network and allows the media serving engine to distribute media content across a
17 network.

18 In another embodiment, the cache policies include policies for distributing
19 media content, handling cache misses, and prefetching media content.

20 In one embodiment, a request for media content is received by a media
21 server from a client. A determination is made as to whether the cache server is
22 functioning as a cache server or an origin server. The request for media content is
23 processed according to a set of cache policies if operating as a cache server. The
24 requested media content is provided to the client if the cache server is operating as
25 an origin server and contains the requested media content.

1

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 illustrates an environment in which a client can request and receive content from a cache server and/or an origin server.

Fig. 2 illustrates an example environment in which multiple clients can request and receive content from a cache server and/or an origin server.

Fig. 3 illustrates another example environment in which data is cached in multiple cache servers.

Fig. 4 is a block diagram showing exemplary components in a cache server.

Fig. 5 is a flow diagram illustrating a procedure for handling a client request to receive pre-recorded content.

Fig. 6 is a flow diagram illustrating a procedure for handling a client request to receive a live broadcast.

Fig. 7 illustrates an example of a suitable operating environment in which the invention may be implemented.

16

17

18

DETAILED DESCRIPTION

The system and methods described herein provide a cache server that performs a variety of cache-related functions. The operation of the cache server is determined by a set of cache policies that are determined by the user of the cache server. Thus, the same cache server is capable of meeting the cache requirements of a variety of users because each user or manufacturer can configure the cache policies to meet their specific needs.

25

1 Fig. 2 illustrates an example environment 200 in which multiple clients
2 212, 214, and 216 can request and receive content from a cache server 204. In this
3 example, each client 212-216 is coupled to cache server 204 via a high-speed LAN
4 connection 218. Cache server 204 is coupled to an origin server 202 via a network
5 206, such as the Internet. Cache server 204 is also coupled to a disk storage
6 device 208, which stores various data, including cached media content (such as
7 media content received from origin server 202). Cache server 204 receives a set
8 of cache policies 210 (also referred to as “policy plug-ins”), which define what
9 functions the cache server performs in certain situations.

10 Origin server 202 stores various pre-recorded media content, such as
11 movies, television programs, concerts, sporting events, and the like. This pre-
12 recorded media content may be distributed to one or more cache servers (such as
13 cache server 204) for redistribution to one or more clients 212-216. Although the
14 distribution of the media content to cache server 204 is via a slow communication
15 link, the cache server 204 is able to redistribute the media content to multiple local
16 clients at a significantly faster rate, thereby providing a higher quality real-time
17 video image to the local clients.

18 Origin server 202 may also receive live broadcast streams, such as news
19 events, press conferences, live sporting events, and the like. These live broadcast
20 streams can be distributed to cache server 204 via network 206. Cache server 204
21 may then redistribute the broadcast streams to one or more clients that have
22 requested the broadcast stream. Thus, a single transmission of the broadcast
23 stream to cache server 204 can satisfy the requests of multiple clients coupled to
24 the cache server, thereby reducing the burden on the origin server 202 and
25 reducing the traffic across network 206.

In the example of Fig. 2, origin server 202 may be a dedicated origin server or may be a cache server (similar to cache server 204) that is capable of functioning as either an origin server or as a cache server. If origin server 202 is a cache server, then it is configured in Fig. 2 as an origin server.

Fig. 2 illustrates a single origin server 202, a single cache server 204, and three clients 212-216. However, it will be appreciated that alternate embodiments may include any number of origin servers, any number of cache servers, and any number of clients, all coupled to one another via one or more data communication networks.

Fig. 3 illustrates another example environment 250 in which data is cached in multiple cache servers 254 and 255. Each cache server 254 and 255 is accessible to different groups of clients. For example, cache server 254 is accessible to clients 262, 264, and 266 via a communication link 268. Although no clients are shown coupled to cache server 255, the cache server may be accessible to any number of clients. Each cache server 254 and 255 is coupled to network 256. An origin server 252 is also coupled to network 256, which allows content to be downloaded from the origin server across the network to the cache servers 254 and 255. Each cache server 254 and 255 is coupled to a disk storage device 258 and a set of cache policies 260.

In this example, the two sets of cache policies 260 are identical such that the operation of the two cache servers 254 and 255 are the same. Each cache server 254 and 255 stores the same cached content from origin server 252 such that a client could be coupled to either cache server and have access to the same cached content. This configuration assists with the sharing of the content distribution burden on both cache servers 254 and 255. If one cache server is

1 reaching its capacity to distribute content, one or more clients may be redirected to
2 the other cache server to balance the content distribution workload between the
3 two cache servers.

4 In an alternate embodiment, the cache servers 254 and 255 do not cache the
5 same content. For example, cache server 254 may cache movie content while
6 cache server 255 caches sports and special event content. Although the two cache
7 servers 254 and 255 may provide the same functions, the cache servers may
8 operate in different manners (e.g., caching different types of content) based on
9 differences in the cache policies applied to the cache servers.

10 Fig. 4 is a block diagram showing exemplary components in a cache server
11 300. Cache server 300 includes a cache engine 302 which performs the various
12 functions necessary to cache media content. Example caching functions include
13 downloading media content from an origin server, distributing media content to
14 one or more clients, redirecting clients to a different server for particular media
15 content, and deleting cached media content to release storage space for caching
16 different media content. A set of cache policies 304 are accessed and utilized by
17 the cache engine 302 in performing its various caching functions. For example,
18 cache policies 304 may define when the cache engine should delete cached media
19 content and when the cache engine should download and cache particular media
20 content for expected future access by one or more clients.

21 In one embodiment, the cache server is a Windows Media Server and the
22 set of cache policies are contained in one or more cache plug-in modules that are
23 loaded in the Windows Media Server and called by the Windows Media Server.
24 Any number of cache plug-in modules can be loaded into the Windows Media
25 Server. The combination of all loaded cache plug-in modules defines the set of

cache policies. In one embodiment, the cache plug-in modules are COM objects, which allows the cache plug-in modules to run in the same space as the Windows Media Server. Alternatively, the cache plug-in modules can run as part of another process on the same computer system or on a different computer system that supports the Windows Media Server. Additional cache plug-in modules can be installed by the system administrator at any time and existing cache plug-in modules can be removed by the system administrator at any time, thereby providing flexible management of the cache policies. Thus, the operation of a cache server can be modified simply by changing the cache plug-in modules in the cache server. Additional details regarding the cache policies are provided below.

Cache server 300 includes various data communication interfaces 306, which allow the cache server to interact with other devices, such as clients, origin servers, and other cache servers. Cache server 300 may use any type of data communication interfaces necessary to communicate with other devices. A media serving engine 308 distributes media content to one or more clients or other servers.

Cache server 300 also includes a data storage device 310 for storing data and other information used by the various components of the cache server. Data storage device 310 works in combination with the disk storage device 208 shown in Fig. 2. The various components shown in cache server 300 may be coupled to one another via one or more data communication busses or other communication links, such as bus 312.

By providing a single cache server 300 that can be configured to operate differently using different cache plug-in modules, system administrators or operators only need to understand how to define cache policies in the cache plug-

1 in modules. Administrators need not understand the internal functioning of the
2 cache server. Cache server 300 combines the necessary components to perform a
3 wide variety of functions that are controlled by the cache policies. Thus, different
4 users or manufacturers can build different cache products using the same cache
5 server 300 platform. The cache engine 302 is a general purpose cache engine
6 capable of performing many different cache functions. The system administrator
7 or operator controls those functions with various cache policies 304, which control
8 the operation of the cache engine 302. Thus, the cache server 300 can be
9 “customized” by each system administrator by defining their own set of cache
10 policies 304.

11 As discussed above, the cache server described herein is capable of
12 performing a set of functions. When and how particular functions are performed
13 is determined by the cache policies. The various functions that a cache server is
14 capable of performing include, for example, downloading content, distributing
15 content, caching content and deleting cached content. The cache policies may
16 determine, for example, how to handle a cache miss, how often to check the
17 “freshness” of cached content, what types of content to cache and when to prefetch
18 content. This flexible cache server architecture allows an administrator to
19 configure a particular cache server to operate in a variety of manners depending on
20 the cache policies associated with the cache server.

21 As mentioned above, the cache server 300 can operate as an origin server
22 and/or a cache server (i.e., caching content on behalf of other servers). If cache
23 server 300 receives a request for content stored in the cache server, the cache
24 server operates as an origin server by providing the requested content. In this
25 situation, the cache server 300 ignores cache policies because the cache server is

1 operating as an origin server. If cache server 300 receives a request for content
2 stored on a different server, the cache server operates as a cache server by caching
3 content from one or more different servers. In this situation, the cache server 300
4 applies the various cache policies to process the cached data in the appropriate
5 manner (i.e., the manner defined by the cache policies).

6 Fig. 5 is a flow diagram illustrating a procedure 500 for handling a client
7 request to receive pre-recorded content. The procedure 500 illustrates the
8 functions performed by the cache server as well as the decisions made by the
9 cache policies. A pair of broken lines 501A and 501B identify the sections of
10 procedure 500 that represent functions performed by the cache server (the sections
11 above line 501A and below line 501B) and the sections of procedure 500 that
12 represent decisions made by the cache policies (the section between lines 501A
13 and 501B). A client establishes a connection to the cache server and requests pre-
14 recorded content from the cache server (block 502). The cache server identifies
15 the policies related to requests for pre-recorded content (block 504). The cache
16 server then identifies whether the requested content is stored on the cache server
17 (block 505). For example, the cache server may examine a URL associated with
18 the requested content to determine whether the content is available on the cache
19 server or whether the requested content is stored on a different server.

20 At block 506, the cache policies determine whether the requested content is
21 acceptable (i.e., not blocked due to adult content or other restricted content). If the
22 requested content requires a verification, the procedure requests a name and
23 password from the user for authentication. If the content is blocked or the client
24 does not provide the proper name and password, the policy terminates the
25 procedure without providing any content to the client.

If the content is acceptable at block 506 (or the proper name and password are provided by the client), then the cache policies determine whether the requested content is available on the cache server (block 508). If the requested content is available, then the cache policies determine whether the cached version of the requested content is an acceptable version (block 510). An acceptable version is determined based on the cache policies (e.g., whether the cached version is the same as the origin server's version or is within thirty minutes of the origin server's version). If the cached version is acceptable, the procedure continues to block 512, where the cache server provides the cached version of the requested content to the client. If the cached version of the requested content is not acceptable, then the procedure branches from block 510 to block 514, where the cache server downloads the current version of the requested content from the origin server. In this situation, the client may experience a slight delay in receiving the requested content as the current version is downloaded. However, the client does not necessarily need to wait until the entire current version is downloaded. Instead, the cache server may begin sending the current version of the content as it is received by the cache server (although the image quality may be reduced due to the typically slow communication links between the cache server and the origin server). As the current version is downloaded (or after it has finished downloading), the cache server provides the current version to the client (block 512).

If block 508 determines that the requested content is not available on the cache server, then the cache policies determine whether to download the requested content or redirect the client to a different server (block 516). If the cache policies dictate that the requested content be downloaded, then the procedure branches to

1 block 514, where the cache server downloads the current version of the requested
2 content from the origin server. If the cache policies dictate that the client be
3 redirected to a different server, then the procedure branches to block 518, where
4 the cache server redirects the client to another server having the requested content.
5 The server to which the client is redirected may be another cache server or an
6 origin server. After redirecting the client to another server to retrieve the requested
7 content, the cache server may (depending on the cache server's cache policies)
8 download the previously requested content so that the cache server has a cached
9 version of the content for future client requests.

10 Fig. 6 is a flow diagram illustrating a procedure 600 for handling a client
11 request to receive a live broadcast. The procedure 600 illustrates the functions
12 performed by the cache server as well as the decisions made by the cache policies.
13 A pair of broken lines 601A and 601B identify the sections of procedure 600 that
14 are represent functions performed by the cache server (the sections above line
15 601A and below line 601B) and the sections of procedure 600 that represent
16 decisions made by the cache policies (the section between lines 601A and 601B).

17 A client establishes a connection to the cache server and requests a live
18 broadcast from the cache server (block 602). The cache server identifies the
19 policies related to requests for live content (block 604).

20 The cache policies then determine whether the cache server should proxy
21 the requested live broadcast (block 608). The cache policy may determine not to
22 proxy the requested live broadcast because the bandwidth quotas for the required
23 communication link have been reached.

24 If the cache policies determine that the cache server will proxy the live
25 broadcast, then the procedure continues to block 610, where the cache server

1 connects to the origin server and establishes a connection to receive the live
2 broadcast. Next, the cache server provides the live broadcast to the client (block
3 612). Any future clients that connect to the cache server can receive the same live
4 broadcast (block 614) until the live broadcast ends (block 616). In a particular
5 embodiment, the cache server stores (i.e., caches) a copy of the live broadcast for
6 viewing by one or more clients after the live broadcast has ended.

7 If the cache server will not proxy the live broadcast, then the procedure
8 branches from block 608 to block 618, where the cache server redirects the client
9 to the origin server. The client then receives the live broadcast from the origin
10 server (block 620) until the live broadcast ends (block 616).

11 The cache policies are typically defined based on an administrator's
12 knowledge of the origin servers, cache servers, clients, networks and other
13 components in the environment in which the cache server is located.
14 Alternatively, the cache policies may have been set by a manufacturer that
15 customizes the cache server and sells the cache server to an end user.

16 The cache policies may also be based on the administrator's knowledge
17 and/or expectations with respect to, for example, the type of content requests
18 received by the cache server, the volume of requests and the timing of the requests.

19 Other decisions the cache policies may make include the types of content
20 that should be cached on the cache server. For example, the cache policy may
21 dictate that the cache server can only cache content from specific origin servers
22 (e.g., origin servers owned by a particular company or organization). Additionally,
23 the cache policy may state that the cache server can only cache small pieces of
24 content (that are inexpensive to copy) or only large pieces of content. The cache

1 policy may also limit the total amount of cached content that the cache server can
2 store at any particular time.

3 Other cache policies determine the frequency with which the cache server
4 checks the freshness of the content stored on the cache server. These policies may
5 also determine whether the cache server is required to have the latest version of
6 the content or whether a “stale” version of the content is sufficient in particular
7 situations.

8 Another group of cache policies handle cache misses (i.e., the cache server
9 does not contain the requested content). When a cache miss occurs, the cache
10 policy determines how to deliver the requested content to the client. The policy
11 may cause the client to wait while the cache server downloads a copy of the
12 requested content, the client may be redirected to a different cache server, or the
13 client may be redirected to the origin server to retrieve the requested content.
14 When a cache miss occurs, the cache policies will determine whether a copy of the
15 requested content should be cached by the cache server to satisfy future requests
16 for the same content by other clients. If the cache policy determines that the
17 content should be cached by the cache server, then another cache policy may
18 determine whether any of the existing cached content should be deleted to release
19 storage space for the new content. If existing cached content is to be deleted, the
20 cache policy will determine which content should be deleted (e.g., the oldest
21 cached content or the least frequently requested content).

22 Another set of cache policies determine when to prefetch content.
23 Prefetching content refers to the process of downloading content from an origin
24 server even though a request for the content has not yet been received from a
25 client. For example, a cache server may download a new movie release or the

1 latest news programs from 3:00 a.m. to 6:00 a.m. in anticipation of client requests
2 for the content later in the day. The prefetching policies determine, for example,
3 the content to download from the origin server and the time period during which
4 the content should be downloaded. If necessary, the cache policies determine
5 which content to delete to make room for the prefetched content.

6 Fig. 7 illustrates an example of a suitable operating environment in which
7 the invention may be implemented. The illustrated operating environment is only
8 one example of a suitable operating environment and is not intended to suggest
9 any limitation as to the scope of use or functionality of the invention. Other well
10 known computing systems, environments, and/or configurations that may be
11 suitable for use with the invention include, but are not limited to, personal
12 computers, server computers, hand-held or laptop devices, multiprocessor systems,
13 microprocessor-based systems, programmable consumer electronics, gaming
14 consoles, cellular telephones, network PCs, minicomputers, mainframe computers,
15 distributed computing environments that include any of the above systems or
16 devices, and the like.

17 Fig. 7 shows a general example of a computer 642 that can be used in
18 accordance with the invention. Computer 642 is shown as an example of a
19 computer that can perform the functions of client computer 212-216, cache server
20 204, or origin server 202 of Fig. 2. Computer 642 includes one or more processors
21 or processing units 644, a system memory 646, and a bus 648 that couples various
22 system components including the system memory 646 to processors 644.

23 The bus 648 represents one or more of any of several types of bus
24 structures, including a memory bus or memory controller, a peripheral bus, an
25 accelerated graphics port, and a processor or local bus using any of a variety of

1 bus architectures. The system memory 646 includes read only memory (ROM)
2 650 and random access memory (RAM) 652. A basic input/output system (BIOS)
3 654, containing the basic routines that help to transfer information between
4 elements within computer 642, such as during start-up, is stored in ROM 650.
5 Computer 642 further includes a hard disk drive 656 for reading from and writing
6 to a hard disk, not shown, connected to bus 648 via a hard disk drive interface 657
7 (e.g., a SCSI, ATA, or other type of interface); a magnetic disk drive 658 for
8 reading from and writing to a removable magnetic disk 660, connected to bus 648
9 via a magnetic disk drive interface 661; and an optical disk drive 662 for reading
10 from and/or writing to a removable optical disk 664 such as a CD ROM, DVD, or
11 other optical media, connected to bus 648 via an optical drive interface 665. The
12 drives and their associated computer-readable media provide nonvolatile storage
13 of computer readable instructions, data structures, program modules and other data
14 for computer 642. Although the exemplary environment described herein employs
15 a hard disk, a removable magnetic disk 660 and a removable optical disk 664, it
16 will be appreciated by those skilled in the art that other types of computer readable
17 media which can store data that is accessible by a computer, such as magnetic
18 cassettes, flash memory cards, random access memories (RAMs), read only
19 memories (ROM), and the like, may also be used in the exemplary operating
20 environment.

21 A number of program modules may be stored on the hard disk, magnetic
22 disk 660, optical disk 664, ROM 650, or RAM 652, including an operating system
23 670, one or more application programs 672, other program modules 674, and
24 program data 676. A user may enter commands and information into computer
25 642 through input devices such as keyboard 678 and pointing device 680. Other

1 input devices (not shown) may include a microphone, joystick, game pad, satellite
2 dish, scanner, or the like. These and other input devices are connected to the
3 processing unit 644 through an interface 668 that is coupled to the system bus
4 (e.g., a serial port interface, a parallel port interface, a universal serial bus (USB)
5 interface, etc.). A monitor 684 or other type of display device is also connected to
6 the system bus 648 via an interface, such as a video adapter 686. In addition to the
7 monitor, personal computers typically include other peripheral output devices (not
8 shown) such as speakers and printers.

9 Computer 642 operates in a networked environment using logical
10 connections to one or more remote computers, such as a remote computer 688.
11 The remote computer 688 may be another personal computer, a server, a router, a
12 network PC, a peer device or other common network node, and typically includes
13 many or all of the elements described above relative to computer 642, although
14 only a memory storage device 690 has been illustrated in Fig. 7. The logical
15 connections depicted in Fig. 7 include a local area network (LAN) 692 and a wide
16 area network (WAN) 694. Such networking environments are commonplace in
17 offices, enterprise-wide computer networks, intranets, and the Internet. In certain
18 embodiments of the invention, computer 642 executes an Internet Web browser
19 program (which may optionally be integrated into the operating system 670) such
20 as the “Internet Explorer” Web browser manufactured and distributed by Microsoft
21 Corporation of Redmond, Washington.

22 When used in a LAN networking environment, computer 642 is connected
23 to the local network 692 through a network interface or adapter 696. When used
24 in a WAN networking environment, computer 642 typically includes a modem 698
25 or other means for establishing communications over the wide area network 694,

such as the Internet. The modem 698, which may be internal or external, is connected to the system bus 648 via a serial port interface 668. In a networked environment, program modules depicted relative to the personal computer 642, or portions thereof, may be stored in the remote memory storage device. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

Computer 642 typically includes at least some form of computer readable media. Computer readable media can be any available media that can be accessed by computer 642. By way of example, and not limitation, computer readable media may comprise computer storage media and communication media. Computer storage media includes volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules or other data. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other media which can be used to store the desired information and which can be accessed by computer 642. Communication media typically embodies computer readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term "modulated data signal" means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as wired network or direct-wired connection,

1 and wireless media such as acoustic, RF, infrared and other wireless media.
2 Combinations of any of the above should also be included within the scope of
3 computer readable media.

4 The invention has been described in part in the general context of
5 computer-executable instructions, such as program modules, executed by one or
6 more computers or other devices. Generally, program modules include routines,
7 programs, objects, components, data structures, etc. that perform particular tasks
8 or implement particular abstract data types. Typically the functionality of the
9 program modules may be combined or distributed as desired in various
10 embodiments.

11 For purposes of illustration, programs and other executable program
12 components such as the operating system are illustrated herein as discrete blocks,
13 although it is recognized that such programs and components reside at various
14 times in different storage components of the computer, and are executed by the
15 data processor(s) of the computer.

16 Thus, a system and method has been described that provide a flexible cache
17 server architecture that is capable of functioning in different manners depending
18 on the set of cache policies that are applied to the cache server.

19 Although the description above uses language that is specific to structural
20 features and/or methodological acts, it is to be understood that the invention
21 defined in the appended claims is not limited to the specific features or acts
22 described. Rather, the specific features and acts are disclosed as exemplary forms
23 of implementing the invention.